www.pipelinepub.com
Volume 19, Issue 11

Taming the AI Wild West

By: Scott St. John, Pipeline

In 1848, James W. Marshall allegedly exclaimed, "there is gold in 'dem hills!" when he spotted precious flakes of the material as he worked on the water wheel of John Sutter's Mill on the American River in what is now Sacramento, California. The discovery sparked what is known today as The California Gold Rush, comprising an estimated 300,000 "forty-niners" that flocked to the West Coast to strike it rich between 1848 and 1855.

Yet, despite the widespread excitement, very few actually struck it rich, or even found gold at all. Instead, most treasure hunters endured major hardships, and countless others died along the way. And, as the masses flooded into the ill-prepared region, lawlessness ensued until proper social structures, governance, and regulations were established. In fact, Sutter was bankrupt by 1853 and the famed American River flooded the city of Sacramento for three months in 1862, in large part due to the debris caused by the practice of hydraulic mining for gold.

Today, there is a similar rush to generative Artificial Intelligence (AI) after ChatGPT extolled the promise of gold earlier this year. But much like the California Gold Rush, this AI frenzy lacks the crucial guardrails needed to navigate the hazards that lie ahead. For end-consumers simply playing with the large-language model, the risk is relatively low. But for enterprises that need to protect their company, customers, employees, shareholders, and brand, the stakes are much higher.

When businesses use generative AI tools for real-world enterprise use cases, they need to be certain that the data set is pure, and the output is accurate and unbiased. To do this, businesses need an enterprise-grade generative AI model that goes beyond public data built for the masses, such as ChatGPT. Businesses require an enterprise-grade generative AI technology that incorporates a diverse range of specific data sources and models that have undergone rigorous filtering, quality controls, and certification, while generating content that is both reliable and ethical. For example, generative AI tools that avoid distributing product content that slants in favor of a specific gender, decisioning that reinforces a racial stereotype, a company spokesperson from promoting conspiracy theories, or a virtual assistant from providing the wrong medical advice to a patient.

In addition to the liability risks, jumping on the generative AI bandwagon with the wrong model can be costly. For example, ChatGPT is an AI-generative engine that automatically generates text based on written prompts in a fashion that is very advanced, creative, and conversational in nature. AI research lab OpenAI launched ChatGPT for public use on November 30, 2022. GPT stands for Generative Pre-trained Transformer and refers to a family of neural network-based large language models (LLMs) developed by OpenAI. ChatGPT is now

one of the largest language models created to date, with a huge neural network that powers the model at 175 billion parameters—consuming massive amounts of computational resources and energy.

It's not ideal for businesses that are being held accountable for ESG goals. From a cost perspective, if a company as large as DoorDash were to replace its prediction engine with this type of generative AI, it could cost them as much as $600 million perday in token consumption, or a staggering $218B per year. The cost model may be too prohibitive for smaller companies to adopt generative AI, causing them to fall back on other technologies—such as Robotics Process Automation (RPA)—or even fallible human workers. And, just as the hundreds of thousands of forty-niners came to realize, promise without protocols can become a real problem. Emerging regulations in the U.S. and U.K. are gradually taking shape to address the safety and social concerns associated with AI, striving to strike a balance between innovation, risks and ethical considerations. As businesses integrate generative AI into their operations, understanding and adhering to emerging regulations will become integral to cultivating a trustworthy and sustainable AI ecosystem. But adherence to the law will not be enough. Legislation is likely to lag behind AI advancements and businesses that want to protect their companies will need to commit to their own high ethical standards to eliminate the risk of bias, inequality, and misguided decisions.

Before generative AI can be successfully adopted by enterprises on a massive scale, the liability, cost, regulatory risk, and sustainability aspects must all be carefully considered. *Pipeline* recently had the opportunity to meet with Madison Gooch, director of Watson and AI Solutions at IBM to discuss these consideration and how generative AI is being used today.

Capitalizing on Enterprise-grade Generative AI

How quickly businesses can capitalize on generative AI depends not only on the projected ROI, but also on competitive factors—such as productivity—that will accelerate the adoption for enterprises and their stakeholders. According to IBM, businesses are already seeing legitimate returns from enterprise-grade generative AI investments for real-world use cases. For example, enterprises are overlaying conversational interfaces on top of business and policy documentation, to improve productivity across customer, customer care, and employee interactions. Summarization is also being widely used by knowledge workers grappling with massive amounts of audio and unstructured data, such as recordings, legal discovery documents, medical transcripts, research documents, or sales meeting notes. Generative AI is being used to both summarize this information and to automatically propagate the relevant bits into other systems. And then there's the generative aspect. Businesses today are using generative AI to help develop a range of content—from legal briefs to educational materials to medical reports—helping workers generate near-final content in a fraction of the time. As businesses see processes streamlined and productivity increased, generative AI can be extended to a wider range of use cases. But this will only be possible by having an enterprise-grade generative AI framework with the appropriate parameters in place.

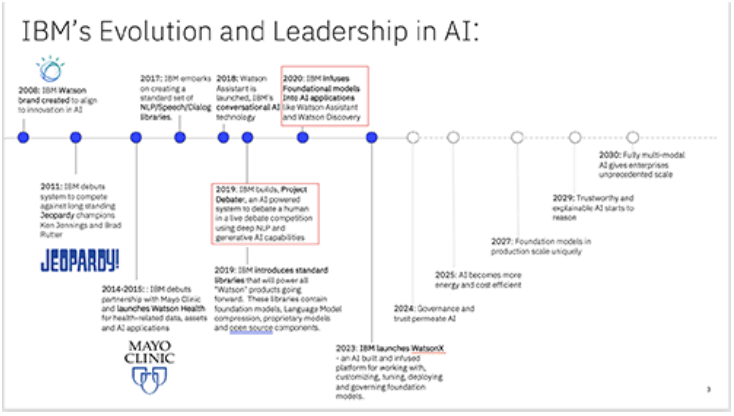

Figure 1 - IBM's AI Evolution
(click here to enlarge)

Picking the Right Generative AI Model

When it comes to addressing governance, choosing the right model before deploying generative AI at scale is key. IBM has been the frontrunner in machine-learning for decades and its investment in AI has been rapidly accelerating in recent years (see Figure 1, above). IBM has been dedicated to creating responsible AI from the very beginning—from the evolution of machine learning to deep learning, dovetailing into AI and now generative AI. IBM's innovation has pioneered the way for the foundation models we recognize as generative AI today. Their contributions have also provided the foundation for their own offerings—including IBM's Watson and watsonx platforms—as well as invaluable contributions to the open-source community.

IBM's AI evolution has been governed by a deep commitment to compliance, governance, and the ethical use of AI. Gooch says this commitment permeates every facet of their approach and underpins IBM's belief that pure input results in ethical output. Using a model that relies on private, proprietary, and vetted third-party training data that undergoes rigorous certification, IBM preserves data integrity and ensures reliable generative AI results. Unlike generative AI driven by large language models that rely upon public or unverifiable sources, IBM's model also provides transparency and auditable models that can trace the lineage of data sources and instillaccountability and trust. Building on that trust is IBM's newest offering, watsonx.governance, the third component in the watsonx platform that also includes watsonx.ai and watsonx.data. Watsonx.governance is designed to help businesses mitigate model risk, detect, manage, and automatically enforce regulatory requirements and address ethical concerns. IBM's generative AI platform allows businesses to direct, manage, and monitor their AI activities and workflows with responsibility, transparency, and auditability. IBM's end-to-end lifecycle approach safeguards against biases, misuse, and misalignment in AI outputs, and its dashboard offers a unique inventory process through "model cards" that are like nutrition labels, where critical information about the model's composition and data source is provided for advanced analytics.

Using foundational models built with responsible AI approaches in mind, including data transparency, businesses can mitigate risk, protect their brand, and expand generative AI workflows for a growing number of tangible, real-world enterprise use cases.

Realizing Generative AI for Enterprise Applications

Unlocking the full potential of generative AI involves not only harnessing its technical capabilities but doing so accurately, ethically, and responsibly. With the proper guardrails in place, businesses can then accelerate their adoption of generative AI workflows, streamline processes, and automate tasks across virtually every functional area—comfortably, confidently, and well beyond what is being done today. By intelligently applying AI and generative AI to do what it does best—such as summarizing large data sets and automating many manual tasks—businesses can repurpose human workers to higher-value and creative tasks.

Realizing the potential of generative AI is possible now with enterprise-grade platforms that leverage responsible models, such as IBM's watsonx. IBM has some great, tangible AI business applications today and will have even more in the near future as new use cases continue to evolve. IBM is currently developing its first set of foundation models for real-world business applications with large language models (LLMs), IT automation models, digital labor models, cybersecurity models, and many more to come.

Enterprises can now begin to deploy generative AI platforms that free up businesses to focus on managing outcomes and expansion versus risk. This makes concepts like digital labor a reality today. IBM is already helping companies remove menial, repetitive, and programmable tasks across marketing, sales, finance, and human resources. Digital employees can communicate, sequence skills on the fly, and put those skills into context by maintaining working memory of past interactions. Workers are able to command, train, and delegate work to digital employees when working on simple tasks and even with more complex decision-making. Finance is relying on digital labor for data analysis, customer care for automated responses, and HR for employee communications. Watsonx is also being used for application modernization and code generation,

transforming the work experience for hundreds of thousands of employees while helping businesses protect their brands. With the proper foundation in place, there's virtually no limit to what generative AI can do, now or in the future.

Charting a Path to Responsible and Reliable AI

While many businesses are rapidly racing to realize generative AI's potential, a rush to adoption without the proper guardrails in place is sure to end with dire, Wild West consequences. Protecting employees, customers, and brands will take a thoughtful approach, with ethics and accuracy an absolute must.

Using a trusted AI model with the ability to meet the needs of each use case will ensure ethical, as well as accurate and reliable, content creation and positive business outcomes. Partnering with a generative AI technology partner such as IBM can help streamline the AI implementation process, as models from reliable collaborators can be harnessed, eliminating the need for businesses to start from scratch. Approaches like those being used by IBM's watson.x platform can provide the foundation and the ability to scale generative AI across enterprise organizations. It can also help usher in a new era of digital labor and workforce enhancement that can unlock true value, accelerate growth, and go a long way toward taming the AI Wild West.